Errors in a team context

James Reason
Department of Psychology, University of Manchester

Niels Ole Bernsen
Centre of Cognitive Science, Roskilde University

Abstract

Despite the large number of laboratory studies of group performance, we still know remarkably little about "team-added" effects on human error. The paper is in three sections. The first summarises the principal findings of laboratory group performance studies (Reason). The second section reviews a number of well-documented accidents with a view to identifying some of the reccurrent patterns of team errors (Reason). The third section presents a principled basis for categorising "team-added" errors (Bernsen).

Introduction

While most analyses of the processes involved in the production of errors have focused upon the single individual acting alone, it is likely that the majority of errors committed in real work environments occur in a team or group context. This raises the important question of how the presence of co-performing team members affects the occurrence, detection and recovery of errors by any given individual.

For the purposes of this paper, we will define a *team* as a group of somewhere between 2 and 12 individuals performing a common task, albeit with specialist roles. This definition would include such diverse groups as control room operators in process plants, aircrew on a flight deck, nurses and doctors working in operating theatres or in special care units, troop command units, workers on a drilling platform and tank or gun crews.

Since the study of group effects upon performance extends back to the 1890s, it is not unreasonable to expect that the social psychological literature of the past 90 odd years would provide some answers to the question posed above. Unfortunately, this is not the case. In the first place, hardly any of the studies have taken account of the *types* of errors elicited, that is, whether they were slips or mistakes. Second, none of them has investigated the effects of co-acting team members upon the likelihood of error detection and recovery. Third, the results of these laboratory group performance studies do not always generalise to real-life work situations.

Section 1: Findings from the social psychological literature

(a) *Audience and coaction effects*

These studies have examined the effects of two kinds of manipulation upon performance: *audience effects*, where the actions are carried in the presence of other people who do not take part in the activity; and *coaction effects*, where the task is carried out with co-performers.

The outcome of these manipulations upon performance is far from clear cut, as the results summarised in Table 1 reveal. Sometimes performance is enhanced and sometimes it is degraded by the presence of passive watchers or coacting others.

A useful attempt to make some sense of this confusing set of findings was made by Zajonc (1965, 1966). He noted that the performance of well-learned and practised activities (repetitive movements, well-learned pursuit rotor tracking, easy mental arithmetic and the like) is generally improved by the presence of an audience, while novel or knowledge-based activities are degraded. He suggested that the primary effect of an audience or co-acting others is to facilitate performance but to impair learning.

Zajonc elaborated this hypothesis further by arguing that the effect of other people is to increase motivation and arousal, which in turn acts to promote the emission of dominant, well-learned responses.

An alternative view has been offered by Baron, Moore and Saunders (1978) and by Freedman, Sears and Carlsmith (1978). This hypothesis focuses upon the distracting effects of others. It is predicted that audience and coaction effects are only likely to occur when people are in a state of conflict between concentrating on the task in hand and attending to task-irrelevant distraction. This idea has found some experimental support (e.g., Baron, Moore & Saunders, 1978).

It is worth noting, in conclusion, that these two hypotheses are not mutually incompatible. Moreover, when taken together, they provide grounds for predicting that audience and coaction influences would render individuals more liable to slips and lapses. The necessary conditions for these actions-not-as-planned are (a) attentional "capture" by distraction or preoccupation, (b) the largely automatic performance of some routine activity, and (c) a need to deviate from routine to accommodate changes in either the plan or the circumstances (Reason & Mycielska, 1982). The outcome is often a "strong-but-wrong" error in which thought, speech or action follows the familiar rather than the intended route.

Table 1. Summarising the results of some early studies on audience and coaction effects.

Investigators	Task	Manipulation	Outcome
Triplett (1897)	Winding a reel	Coaction	Improved performance
Meumann (1904)	Ergograph	Audience effect	Improved performance
Allport (1920)	Word associations, cancelling vowels in newsprint, simple multiplications	Coaction	Improved performance
Allport (1920)	Refuting false syllogisms	Coaction	Degraded performance
Travis (1925)	Pursuit rotor	Audience effect	Improved performance
Dashiell (1930, 1935)	Multiplication problems, analogies, word association	Audience and coaction effects	Worked faster but with reduced accuracy
Husband (1931)	Finger maze	Audience effect	Interference with learning
Pessin & Husband (1933)	Finger maze	Audience effect	Interference with learning
Pessin (1933)	Learning nonsense syllables	Audience effect	Interference with learning
Bergum & Lehr (1963)	Vigilance task	Audience effect	Detection accuracy 34% higher than when alone

(b) *The group polarisation phenomenon*

In 1957, Ziller made the seemingly counter-intuitive observation that group decisions are often riskier than those taken by its individual members when working alone. This so-called *risky shift effect* was rediscovered independently by Stoner (1961), using the

Choice Dilemmas Questionnaire (CDQ). The CDQ outlines 12 hypothetical situations in which the subject has to choose between a safe course of action and a more attractive, but riskier one. The dilemmas included such situations as choosing between a moderately satisfactory job and more attractive one in a company that is financially precarious, or between playing for a safe draw in chess match or playing for a win and thus risking defeat. Other situations involved making a safe domestic investment or a riskier but more potentially profitable one abroad, or deciding whether or not to marry. The CDQ became the favoured research tool in the many later studies that sought to replicate these findings (see Myers, 1973).

These studies confirmed that risky shifts occurred reliably on most of the CDQ items. But cautious shifts were found on certain items, particularly the investment and marriage ones.

Many explanations have been advanced -- familiarization, leadership effects, diffusion of responsibility -- but only one, Brown's (1965) *value hypothesis*, is held to provide an adequate explanation of both risky and cautious shifts (known collectively as *group polarization*).

The argument goes like this: In Western cultures, we tend to admire risk-takers more than those who are timid or cautious. We like to consider ourselves at least as willing to take risks as other people. When completing the CDQ on an individual basis, subjects assume (in the absence of any evidence to the contrary) that they are making riskier choices than they actually are. But in the group situation, many people discover that there are others with higher risk levels. In order to restore their self-images, they change their decisions in the direction of greater risk-taking.

This accounts for the risky shift. The cautious shift can be explained by a minor elaboration. While risk-taking is admired in general, there are certain situations (investments, marriage, etc.) where it is considered foolhardy. In these high-consequence circumstances, caution rather than risk-taking is the more valued option. Once discrepancies are discovered in the group situation, decisions on these items shift in the more cautious direction.

There seems little doubt that interpersonal comparisons of this kind can shift group decision-making away from the norm of individual decisions, but whether these social influences render a team more liable to make mistakes has not been clearly demonstrated. It could work either way, depending upon the situation. Given the almost infinite range of circumstances under which real-life decisions are taken, it is doubtful whether the group polarisation phenomenon could be of much help in predicting either the occurrence or the nature of future mistakes, but it might help to cast some light upon what led to some of the more impenetrable mistakes of the past.

(c) *Groupthink*

This phenomenon has been too widely discussed elsewhere to require much coverage here. Whereas Cyert and March (1963) observed that errors in a collective setting tend to cancel one another out, or what they termed *bias discounting*, Janis (1972, p.7) asserted that "...all the well-known errors stemming from the limitations of an individual and of a large organization can be greatly augmented by group processes that produced shared miscalculations."

The "groupthink syndrome" arises in small elite groups and is characterised, among other things, by a disregard of prior warnings of likely failure, an unswerving belief in the group's own rightness, a shared illusion of unanimity and a strong reluctance on the part of any group member to express doubts or to disrupt the group's unanimity in any way. However, should such doubts be expressed, strong pressure is put upon the deviant member to bring him or her back to the agreed party line.

Janis's analyses of planning teams are among the few that specifically consider the effects of group membership upon the *quality* of the resulting decisions. As Janis makes clear, however, the same group (i.e., Kennedy's advisers) acting under the same social pressures can make highly successful decisions (i.e., the handling of the Cuban missile crisis) as well as disastrous ones (i.e., the Bay of Pigs landing). Even "groupthink" is not a sufficient condition for the making of mistakes.

(d) *Conformity effects*

Perhaps the best known of all conformity studies, that by Asch (1951), demonstrated very clearly the profound effects that group pressures may have on the judgement, even the perceptions, of individuals. The experimental technique was very simple. Groups of eight subjects (7 stooges and one unsuspecting subject) were required to match the length of a given line with one of three unequal lines. The judgements were given publicly. In the midst of this otherwise tedious procedure the subject suddenly found his judgement contradicted by the entire group and these contradictions were repeated several times during the course of the experiment. The errors of the majority were large (ranging between a half and one-and-threequarter inches). Under such conditions, a substantial minority of subjects yielded to the majority pressure.

Asch identified three types of yielders: (a) those who unconsciously distorted their perceptions and did not realise that the judgements of the majority were incorrect; (b) those who consciously decide, after some agonising, that their perceptions were inaccurate and succumb to the strong pressure to join in with the majority view; and (c) those who neither doubt the evidence of their senses nor believe that they are wrong, but simply appear to "give in" to the group pressure, either because they have little investment in the relative lengths of lines, or because they do not wish to feel different from or inferior to the majority.

These individual differences raise some interesting issues for the student of human error. The first group, the *perceptual distorters*, were profoundly in error. They failed to

interpret the evidence of their senses correctly. The second group, the *judgemental distorters,* preferred error to non-conformity. Some of the third group may not have cared much one way or the other and simply opted for the easy way out. These people were probably not in error, though those that yielded out of social anxiety may well have been and certainly would have felt so when the situation was eventually explained to them.

These and numerous other studies upon the effects of group pressure indicate that the desire to conform can provoke errors of various kinds. Processes similar to those described by Asch probably contributed to the *groupthink phenomenon*, particularly when individual group members may well have felt flattered to have been a part of such elite groups. But it must be remembered that the majority of Asch's subjects elected not to yield to group pressure. Clearly then, group pressures, though powerful, do not guarantee the commission of error.

Section 2: Case study evidence

When laboratory studies fail to provide convincing answers, it is necessary to look elsewhere. Accident and incident case studies, while selective and incomplete, do provide a reasonable account of the forms of team errors. In this section, we will summarise the team aspects of a small number of well-documented accidents and near-misses with a view to identifying some of the recurrent varieties of team-related errors.

(a) Independent but contingent errors.

At Paragon Station Hull, in 1929, two signalmen each made an independent error which, together, created a tiny chink in the station's highly sophisticated (for those times) defences. The outcome was that the outgoing Scarborough train was switched onto a track upon which the Withernsea train was approaching. A number of passengers died and several were injured in the subsequent collision. The first error was that of Signalman Gibson who set the signal to danger prematurely, thus releasing the points from the first safety lock. This created a brief "window of fatal error opportunity" within which Signalman Clark erroneously pulled points lever number 95 instead of 96. This was all that was required to switch the Scarborough train onto the Withernsea line.
The pattern here then was of two independent errors, occurring very closely in time and penetrating the railway system's elaborate anti-collision defences.

A similar pattern of errors, but on a larger scale, contributed to the Tenerife runway disaster in 1977. The Pan American crew taxied past Exit 3, the turn-off directed by the air traffic controller. The KLM captain started his takeoff run before receiving takeoff clearance. The KLM co-pilot allowed the takeoff to proceed, even though he was

aware that an error had been committed. Together, the two sets of errors were sufficient to bring the two jumbos into collision on the runway.

There was also an added factor: the social psychology of the cockpit. The KLM captain was the airline's chief training pilot and an extremely senior figure. It seems reasonable to suppose that the very much more junior co-pilot was hesitant to call too much attention to the captain's precipitate takeoff.

(b) Shared errors.

Another common pattern of team error is where all members of the group make the same misdiagnosis or engage in the same inappropriate actions. These shared errors were evident at Three Mile Island operating crew and in the Eastern Airlines Everglades crash.

(*c) Sequential errors.*

Yet another pattern involves a sequence of independent errors, each one creating the conditions for the next. At an incident at the Oyster Creek Boiling Water Reactor in 1979 a junior operator made an action slip, closing all four pump discharge valves instead of the two intended. This caused the crew to misinterpret the level of water in the annulus, despite a triple low-level alarm. As a result they applied the wrong corrective actions (cycling) to the isolation valve. Similar sequences of errors contributed to the Charge of the Light Brigade (1854), the Ginna nuclear power station incident of 1979 and the Hillsborough football stadium disaster of 1988.

(d) Single uncorrected error(s).

In this pattern, there is often only one significant active failure, usually committed by a senior member of the team. The problem lies in the failure of the remainder of the crew to detect and correct the error. This pattern has contributed to a large number of aircraft accidents. A critical factor appears to be the unwillingness on the part of other crew members to challenge the authority of the captain (e.g., the Papa India crash of 1972).

(a) Independent but contingent errors

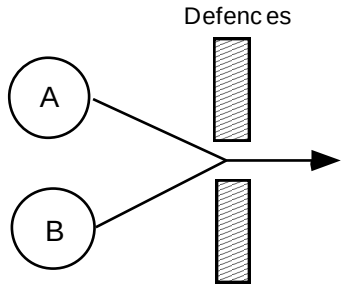

hared errors

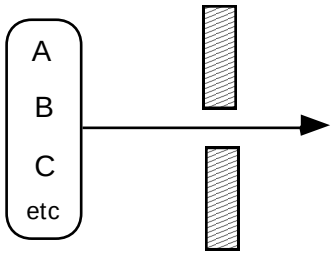

(c) Sequential errors

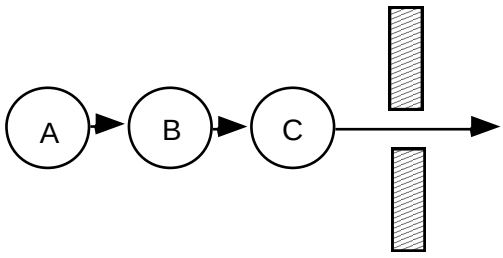

(d) Single uncorrected error

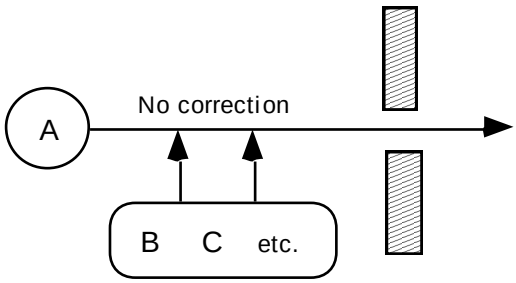

Figure 1. Four patterns of team error

On the basis of these case study analyses, it is possible to identify a number of "team-added" error factors. These are listed below:

Presumption of others' competence
Unawareness of others' errors
Distraction
Diffusion of responsibility
Self-image (self-esteem) preservation
Authority problems
Collective "cognitive lockup"
Personality problems, etc. etc.

Section 3: Towards a taxonomy for team-added errors

The factors listed above would seem to represent an important fraction of possible team-added errors. However, the literature on errors in team contexts does not appear to contain higher-level structuring principles which could help us get a theoretical handle on this domain of human error.

The unresolved questions are: (1) How do we structure this particular error domain in order to understand why errors of, i.a., the above types seem to be prominent in team work? (2) Why are these types of error characterizable as *team-added* errors? And (3) is it possible to propose a taxonomy of team errors which might allow us to generate, top-down and in a principled manner, all or most generic types of team-added error? In what follows, I will propose a taxonomy which helps answering questions (1) to (3).

First of all, what is team work? It seems clear that team work at least comprises all kinds of work organisation in which several people work together in order to accomplish results that they could not accomplish as efficiently, or as quickly, or as safely, or as reliably, or not at all, if they did not work together. However, underneath any such loose and general characterization of team work many different varieties of team work can be found including, e.g., the mere division of labour without any substantial team member interaction, and it is not the purpose of this paper to attempt a comprehensive taxonomy of these. Instead, we shall take a specific kind of team work as our guiding example. This is the *closely-knit team work* one observes in relatively small, closely interconnected groups with a high degree of mutual dependency, such as cockpit airline crews, space crews, tank crews in combat, emergency surgery teams, or teams of mountain climbers. The reasons for starting by looking at such teams is not only that these were the teams studied by Reason but, more fundamentally, and that was probably his reason for studying them, that such teams would seem to be the prototypical teams in our common cognitive model of teams engaged in team work. This means that all or most other kinds of team work

could be (indeed, if I am correct about this, *are*) conceived of as different varieties of team work derived from the core conception. This principle of conceptual organisation, in which a complex cognitive model is organised around a core or prototypical conception, is ubiquitous in human conceptual systems (Lakoff 1987). If closely-knit team work does indeed form the core of our cognitive model of team work, it is likely that a successful taxonomy of team added errors in closely-knit team work might serve as a core model from which error types in other team work contexts could be derived and understood.

A cognitive model of closely-knit teams.

Let us take a closer look at closely-knit teams. Undoubtedly, such teams fit the various kinds of rationale for having teams rather than independent individuals to accomplish a given task , i.e., they make task accomplishment safer, or quicker, or more reliable, or more efficient, or ..., etc., or they simply make task accomplishment possible in the first place. These kinds of rationale are also valid for teams which are not closely-knit teams. But why do closely-knit teams make such accomplishments possible?

I suggest that we have an *ideal cognitive model* of the closely-knit team. We do not have this ideal cognitive model because we believe, and in that case erroneously, that reality is always like that. We have the model partly because it provides the normative justification for assigning some tasks to closely-knit teams. The ideal cognitive model equips closely-knit teams with a number of characteristics which justify the setting-up and use of closely-knit teams for certain tasks and at the same time explain why they are able to do the task.

The characteristics of closely-knit teams in the ideal cognitive model include the following:

1. The members of a closely-knit team freely communicate, i.e. they communicate with each other in an unconstrained manner.

2. The communication among members of a closely-knit team is totally objective, i.e., the communication does not produce any undesirable side-effects.

3. The communication (verbal and otherwise) among members of a closely-knit team is fully transparent to the participants in communication.

4. The members of a closely-knit team are a able to monitor andcontrol each other's behaviour, i.e., they always act as autonomous team members in the collaboration.

5. All members of a closely-knit team are generally competent in performing their part of the task, i.e., each member normally knows what to do and how to do it correctly.

6. The members of a closely-knit team have a fully transparent distribution of responsibility while doing their part of the task assigned to the team. The members know the responsibilities of each other.

(1) to (6) specify the relationships between the individual members of closely-knit teams. There is a further characteristic which has a different status since it concerns the relationship between the team as a whole and the task:

7. The skills and experience of the members of a closely-knit team are jointly necessary and sufficient to perform the task successfully.

(7) is clearly necessary to justify the use of a specific closely-knit team in a specific task and it does form part of our ideal cognitive model of such teams. The model is a model of teams which are up to the task set to them.

We saw that the ideal cognitive model justifies the use of closely-knit teams in certain circumstances. The characteristics of the model are at the same time normative rules which state how a closely-knit team should ideally behave. An equally fundamental reason why we have this cognitive model is the following. Reality actually does contain such ideal teams, at least for limited periods of time and for specific tasks. So the ideal model is an abstraction from experience. We precisely tend to think, at a certain level of abstraction, of such actual teams and their task performance when discussing closely-knit teams and their deployment. Some central features of such teams typically are that their members have worked together for a long time, they know each other well, and the team is relatively small. Features like these serve to explain why their cooperation manages to satisfy characteristics (1) to (6) above. A typical task set for a closely-knit team is one which requires closely coordinated real-time action under time pressure by the team members. It is also typical of such tasks that team failure may lead to serious losses of human life and property.

Breakdown of the normative presuppositions

If we assume that (1) to (6) specify most of the salient characteristics belonging to the ideal cognitive model of a closely-knit team, then we can start generating error types top-down. For, alas, this is not an ideal world, and in most actual closely-knit teams the ideal presuppositions above break down time and again. This is normal, and in most cases the consequences are not serious. But sometimes they are, as in the cases described in Section X of this paper. Let us assume a closely-knit team of two persons, A and B. Each of these characteristics, which we alternatively describe as *presuppositions*, may be violated independently of the others.

Breakdown of presupposition (1): A and B do not freely communicate. Analysing this negated presupposition, we get two different conditions:

(a) A and B do communicate, but not freely;
(b) A and B do not communicate at all.

1 (a) is the locus for situations in which A does not want to tell B about a mistake made by A, and hence of the error found by Reason due to *self-image preservation* (by A). 1 (a) is also the locus for situations in which A does not want to ask B a question which is crucial to correct task performance, and hence of the error found by Reason due to *authority problems*. In this category we also find errors (not included in Reason's error sample) being committed because of the fact that A and B simply happen not to communicate when they should have done so. We might call such errors *communication slips*.

1 (b) is the locus for situations in which errors are made because A more generally does not want to talk to B or vice versa (or both). Since in such circumstances crucial information may be needed by one of the team members but for the reason stated cannot be obtained, we get the error found by Reason which was due to *personality problems*.

Breakdown of presupposition (2): The communications between A and B are not totally objective, so that their communication produces undesirable side-effects. One such side-effect is *distraction*, also found by Reason, which is produced when the turn of communication towards irrelevant topics leads to lack of attention to the task. Other such possible side-effects include, e.g., the *spread of panic* from one team member to another.

Breakdown of presupposition (3): The communication (verbal or otherwise) among A and B is not fully transparent. Violation of communicative transparency takes many forms which may all lead to serious error: what A communicates to B may be unintentionally unclear, ambiguous, incomplete, misleading, etc., and even if this is not the case, B may not correctly understand the intended message for a variety of reasons. Let us call all errors of this category *errors due to unsuccessful communication*. A very different category of errors due to breakdown of presupposition (3) comprises the errors produced when A *intentionally* misleads B. In such cases, which may be rare given the fact that the team members rely on each other in accomplishing their task, one might expect to find personality problems, jealousy, or other factors behind A's behaviour.

Breakdown of presupposition (4): A and B fail to monitor and control each other's behaviour, so that they do not succeed in acting as autonomous team members in the collaboration. In this case, we get Reason's observed *non-detection of errors* committed by other team members. Clearly, A and B should not always monitor and control each other, since this would seriously compromise the advantages of their acting as a team.

However, in many different specific cases they actually should monitor and control each other. There are, for instance, cases in which A relies on B although he should not do that. If B communicates some surprising observation relevant to the common task, A might have to look for himself and possibly also verify if B's condition is still O.K. Or if B is new to the collaboration, A should not fully rely on B in difficult circumstances until he has seen how B is able to cope with these. This case is related to breakdowns of presupposition (5) below. Or B may be indisposed on a particular day, and A should have noticed this instead of blindly relying on B. It should be noted that A's (unjustified) reliance on B may also be due to a non-transparent distribution of responsibility among the team members (presupposition 6).

Breakdown of presupposition (5): A or B is not generally competent in performing his part of the task. Here we have the errors Reason found due to the *others are competent* assumption. As we have seen, the monitoring and control of the activities of other team members serve as a "backup" here. Since team members cannot (and should not) in practice exert continuous control of each other's actions, and since they cannot test each other's skills and understanding in every possible situation before it arises, A may lack some skill or understanding which becomes noticed by B only when it is too late.

Breakdown of presupposition (6): A and B do not have a fully transparent distribution of responsibility among themselves. In this case, sub-tasks crucial to the performance of the overall task of the team may not be performed at all or may be performed too late. A believes that B has the responsibility for doing sub-task S1, and B believes that this is A's responsibility. The errors Reason found due to *diffusion of responsibility* belong in this category.

We have now dealt with all the categories of error identified by Reason except one, the *collective "lockup"* in a false interpretation of the situation. Where does this error type belong in the taxonomy? It could be produced from several types of breakdown within our ideal cognitive model of closely-knit teams. It could be produced from A's relying too much on B (who might be the team leader). This would constitute a new kind of authority problem. A and B do communicate but due to the job hierarchy, A does not think independently, so that we get a breakdown of presupposition 4. It could be produced from spread of panic (breakdown of presupposition 2). Or the "lockup" could be produced as follows:

Breakdown of presupposition (7): The skills and experience of A and B are not jointly sufficient to the successful performance of the task even though both A and B are, by any standard, sufficiently competent. In an extreme situation, A may rely on B just because B has an idea, any idea, whereas A has none. This is probably a borderline case for applying the concepts of error and team-added error. In other instances, in which it makes sense to say that improvements in training or equipment might have made it possible to avoid a particular accident due to breakdown of presupposition (7), the error may lie with those who thought that the team consisting of A and B and

their equipment would be up to the task set to them. So breakdowns of presupposition (7) do not always count as a team-added errors. Had the team itself undertaken the task in question, we actually would have a "team-added" error, albeit of a particular kind.

Taking stock

It seems that we have a common, ideal cognitive model of closely-knit teams. The model represents a well-working, closely-knit team at some level of abstraction. This means that the model does not specify the exact number of team members, the exact nature of the task and its environment, whether or not there is a team leader, and so on. Rather, the model is constituted by a number of structural characteristics. Thus, the model includes a number of specific and mutually independent presuppositions about the nature of such teams and the way their members work together. The presuppositions form part of the background justification for applying closely-knit teams to particular tasks rather than individuals or teams of other (more loosely coupled) kinds. When one or several of these presuppositions break down in real life team work, as is often the case and mostly without producing any harmful consequences, the result may be a team-added error which causes a more or less serious accident.

Section 2 of this paper described a number of such accidents and identified the particular (closely-knit) team-added error responsible for them. Using the ideal cognitive model, it was possible to define each of these errors as resulting from a particular breakdown in the presuppositions of the ideal cognitive model of closely-knit team work. In addition, when studying such breakdowns we have found a number of team-added errors which were not included in case material presented above.

Thus, the ideal cognitive model has proved useful as a set of generating principles of team-added errors and in providing a taxonomy for such errors. Clearly, we are not able to generate all possible, specific types of team-added error in closely-knit teams. What we have done is to provide a framework in the sense of a number of generic types of team-added error in which all or most of these might fit.

Loosely-coupled teams.

Since the properties of the ideal cognitive model of closely-knit teams are abstract structural characteristics, it becomes possible to suggest that the negation of one or several of them in some cases will produce *different cognitive models* related in various ways to the core model. These models can be seen as derived variations on the ideal model. Negation of each of the structural properties (1) to (7) above produces

continuous scales of degradation, as it were, of the team work. Team member communication may be more or less constrained, or may produce more or less undesirable side-effects, the responsibility distribution among team members may be more or less vague, team member competence may be more or less sufficient to the task, etc. So negation by itself is not sufficient to produce different cognitive models of closely-knit team work. Rather, reality has produced some of these different model already:

One is the *new team* whose members are unknown to each other and which now has to accomplish a task which requires closely-knit team work. This may happen, e.g., during emergencies in all kinds of circumstances and environments. Such a team will have to define itself with respect to the structural dimensions of the ideal cognitive model. Another is the *incompetent team*, i.e., the team whose members are not competent in performing their parts of the task. This team will probably fail. A third variation, the *playing team*, can be seen in many different kinds of game, from soccer and tennis doubles to various card games. Such closely-knit team work is primarily characterized, although the difference can sometimes be difficult to discern in practice, by dealing with a different generic task domain, i.e., that of "play" rather than work. But all such teams are variations of closely-knit teams and can be characterised along the structural dimensions of the ideal cognitive model.

What happens if we leave the domain of closely-knit teams and the corresponding, generic task domain and move into more loosely-coupled cooperative work and its corresponding, generic task domain ? The main task domain difference seems to be that, typically, tasks no longer require closely coordinated real-time action under time pressure by the team members. That is why one can afford using more loosely-coupled team work. This is a vast domain. We may have huge teams, extremely general tasks, such as "running the company", complex hierarchies of command, a wide variety of command and control structures, large amounts of written instructions, rules, and regulations, or long time spans for the measurement of success and failure.

Such differences from closely-knit team work seem to have the following implication. There is no single, core cognitive model of loosely-coupled team work, except in a highly abstract sense: there has to be a task to be solved by the team through the use of some kind of collaboration between team members; team members have to communicate with other team members to some extent; some kind of monitoring and control of the team work is necessary; team member competence is still a normative requirement; and to some extent, at least, transparent distribution of responsibility among team members is normatively required, as is the sufficiency of the team to the task set to it. But these characteristics are *unfocused* in the sense that they fit all kinds of loosely-coupled teams with no preference for any particular subset of them. Incidentally, they also fit closely-coupled teams, since they constitute our abstract concept of team work in general.

To identify more concrete cognitive models of loosely-coupled team work, we have to go down one step in the abstraction hierarchy. There it becomes possible to characterise a particular team in analogical terms as operating "like a family", "like a ship crew", "like the mafia", "like a Byzantine bureaucracy", "like an efficient private company", and so on. If we want to use the conceptual apparatus of cognitive semantics, we might say that the concept of a closely-knit team is a *basic-level concept* having a rich structure and a clearly identifiable class of core instances. By contrast, the concept of a loosely-coupled team is not a basic level concept but a *higher-level, abstract concept*. The basic-level concepts of loosely-coupled teams are to be found one level lower, being concepts or models like "family", "mafia", "ship crew", etc. These concepts have a rich structure and a clearly identifiable class of core instances. If we proceed one level further down, we find the concepts and models of particular types of family, chip crews, companies, public bureaucracies, and so on.

This explains why the closely-knit team could be usefully studied for the purpose of deriving a taxonomy of team-added errors, whereas the concept of loosely-coupled teams cannot fruitfully be studied for this purpose. Possible taxonomies of team-added errors for loosely-coupled teams would have to be developed from a series of basic-level concepts in that domain.

The concept of loosely-coupled team work, then, seems to be meaningful only as contrasted with the concept of closely-knit team work which is a basic-level concept. It is in some sense the "real" and "ideal" team work which becomes diluted when we turn toward loosely coupled team work which is as much created for the purpose of mere division of labour as it is created for the purpose of having team members working together. The expression "*mere*" *division of labour* refers to a task, T, which can restlessly be broken down into sub-tasks t1, t2, and t3, say. The sub-tasks can be performed by different workers working independently of one another. When they have all performed their respective sub-task, T has been accomplished. In this sense, mere division of labout required a minimum of coordination and minimal team member interaction. This leads to one *final hypothesis*. It is that the additional team-added error types to be found in loosely-coupled team work, i.e., types of error which have not already been found in the study of closely-knit team work, are types of error which originate in breakdowns of schemes developed in order to combine characteristics of closely-knit team work with the characteristics of mere division of labour. If that is true, the main approach to specific types of team-added error in loosely-coupled team work will be to look into the *coordination and integration schemes* characterising the different basic-level categories of loosely-coupled team work; define the ideal cognitive model of these; and study the various ways in which they may break down.

Summary and conclusions

A review of the social psychological literature revealed four topics having a bearing on team errors: audience and coaction effects, the group polarisation phenomenon,

"groupthink" and conformity effects. Although useful in identifying group processes leading to error, the laboratory studies revealed little about either error types or error detection. There was also a suspicion that some of these well-established laboratory phenomena did not readily transfer to the real world.

The second part of the paper examined a number of well-documented accident case studies in an attempt to identify some recurring patterns of team error. Four such patterns were considered: independent but contingent errors, shared errors, sequential errors, and uncorrected errors. A number of "team-added" error factors were also noted: presumption of competence, unawareness of others' errors, distraction, diffusion of responsibility, self-image preservation and authority problems.

The third section attempted to provide something that is conspicuously lacking in either the literature or the case study material: namely, a set of higher-level taxonomic principles. This is founded on an ideal model of closely-knit teams. Seven normative characteristics were formulated and the consequences of the breakdown of each considered. The section ended with a consideration of loosely-coupled teams. These were seen to generate additional errors forms related to the coordination and integration of group activities.

References

Allport, F.H. (1920) The influence of the group upon association and thought. *Journal of Experimental Psychology*, 3: 159-182.

Asch, S.E. (1952) *Social Psychology*. Englewood Cliffs, NJ: Prentice-Hall.

Baron, R.S., Moore, O., & Sauders, G.S. (1978) Distraction as a source of drive in social facilitation research. *Journal of Personality & Social Psychology*, 36: 816-214.

Bergum, B.O. & Lehr, D.J. (1963) Effects of authoritarianism on vigilance performance. *Journal of Applied Psychology*, 47: 75-77.

Brown, R. (1965) *Social Psychology*. New York: Free Press.

Cyert, R.M. & March, J.G. (1963) *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall

Dashiell, J.F. (1930) An experimental analysis of some group effects. *Journal of Abnormal and Social Psychology*, 25: 190-199.

Dashiell, J.F. (1935) Experimental studies of the influence of social situations on the behavior of individual human adults. In: C.Murchison (ed). *A Handbook of Social Psychology*. Worcester, Mass: Clark University Press.

Freedman, J.L., Sears, D.O., & Carlsmith, J.M. (1978) *Social Psychology*. Englewood Cliffs, NJ: Prentice-Hall.

Husband, R.W. (1931) Analysis of methods in human maze learning. *Journal of Genetic Psychology*, 39: 258-277.

Janis, I.L. (1972) *Victims of Groupthink*. Boston: Houghton Mifflin

George Lakoff (1987) *Women, Fire, and Dangerous Things*. Chicago: The University of Chicago Press.

Meumann, E. (1904) Haus- und Schularbeit: Experimentie an Kindern der Volkschule. *Die Deutsche Schule*, 8: 278-303, 337-359, 416-431.

Myers, D.G. (1973) Summary and bibliography of experiments on group-induced response shift. *Catalog of Selected Documents in Psychology*, 3: 123.

Pessin, J. (1933) The comparative effects of social and mechanical stimulation on memorizing. *American Journal of Psychology*, 45: 263-270.

Pessin, J., & Husband, R.W. (1933) Effects of social stimulation upon human maze learning. *Journal of Abnormal and Social Psychology*, 28: 148-154.

Reason, J. & Mycielska, K. (1982) *Absent-Minded? The Psychology of Mental Lapses and Everyday Errors.* Englewood Cliffs, NJ: Prentice-Hall.

Stoner, J.A.F. (1961) *A Comparison of Individual and Group Decisions Involving Risk*. Master's Thesis. MIT School of Industrial Management, Cambridge, Mass.

Travis, L.E. (1925) The effects of a small audience upon eye-hand coordination. *Journal of Abnormal and Social Psychology*, 20: 142-146.

Triplett, N. (1897) The dynamogenic factors in pacemaking and competition. *American Journal of Psychology,* 9: 507-533.

Zajonc, R.B. (1965) Social facilitation. *Science*, 149: 269-274.

Zajonc, R.B. (1966) *Social Psychology: An Experimental Approach*. New York: Wadsworth.

Ziller, R.C. (1957) Four techniques of group decision making under uncertainty. *Journal of Applied Psychology*, 41: 384-388.